

7

prompt being provided for presentation to the user via a user interface output device of the client device; receiving additional natural language input provided by the user in response to providing the natural language prompt; determining a value for the parameter based on the additional natural language input; selecting a particular third-party agent of a group of third-party agents that can perform the intended action; transmitting a third-party invocation request that includes the value for the mandatory parameter, where the transmitting is to the particular third-party agent via one or more network interfaces; receiving responsive content from the third-party agent in response to transmitting the intended action and the value, the receiving being via one or more of the network interfaces; and providing output that is based on the responsive content for presentation to the user; receiving further additional natural language input provided by the user in response to providing the output; performing an additional voice to text conversion to convert the further additional natural language input to further text; and transmitting the further text to the particular third-party agent.

In some implementations, a method performed by one or more processors is provided and includes: receiving, via one or more network interfaces, input transmitted by a client device, the input being free-form natural language user input; determining an intended action based on the input; identifying a mandatory parameter that is stored as mandatory for the intended action; determining that the input lacks specification of any value for the mandatory parameter; in response to determining that the initial text lacks specification of any value for the mandatory parameter: generating a prompt that is based on the mandatory parameter and that does not solicit input on an optional parameter stored as optional for the intended action, and transmitting, to the client device, the natural language prompt as a reply to the input; receiving, via one or more of the network interfaces, additional input transmitted by the client device in response to the prompt, the additional input being free-form natural language user input; determining a value for the parameter based on the additional natural language input; determining an additional value for the optional parameter based on the additional natural language input; selecting a particular third-party agent that can perform the intended action based on both the value and the additional value; and transmitting a third-party invocation request that includes both the value and the additional value. The transmitting is to the particular third-party agent via one or more of the network interfaces.

In addition, some implementations include one or more processors of some computing devices, where the one or more processors are operable to execute instructions stored in associated memory, and where the instructions are configured to cause performance of any of the aforementioned methods. Some implementations also include one or more non-transitory computer readable storage media storing computer instructions executable by one or more processors to perform any of the aforementioned methods.

It should be appreciated that all combinations of the foregoing concepts and additional concepts described in greater detail herein are contemplated as being part of the subject matter disclosed herein. For example, all combinations of claimed subject matter appearing at the end of this disclosure are contemplated as being part of the subject matter disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of an example environment in which implementations disclosed herein may be implemented.

8

FIG. 2A is a flowchart illustrating an example method according to implementations disclosed herein.

FIG. 2B is a flowchart illustrating some implementations of the blocks of the flowchart of FIG. 2A in more detail.

FIG. 3 is a flowchart illustrating an example method that may be performed by a third-party agent according to implementations disclosed herein.

FIGS. 4, 5, and 6 each illustrate a user, a client device, and an example dialog between the user, an automated assistant associated with the client device, and a third-party agent, according to implementations disclosed herein.

FIG. 7 illustrates an additional client device, and an example dialog between a user of the client device, an automated assistant associated with the client device, and a third-party agent, according to implementations disclosed herein.

FIG. 8 illustrates an example architecture of a computing device.

DETAILED DESCRIPTION

In some situations, in order to invoke a particular third-party (3P) agent via an automated assistant, a user must provide input that explicitly invokes that particular 3P agent. For example, to invoke a “movie ticket purchase” 3P agent named “Hypothetical Agent”, the user must know to speak a “hot word/phrase” for the 3P agent such as “Order tickets with Hypothetical Agent”. Such explicit invocations require the user to know at the outset which 3P agent is most appropriate for an intended action of the user, and sends the user directly to that 3P agent for attempted resolution of the intended action via interaction with the 3P agent.

However, it may often be the case that the 3P agent directly invoked by the user isn’t able to perform the intended action in a manner desired by the user. This may waste user and computational resources as the user must first interact with the 3P agent, determine the 3P agent is not appropriate, then attempt to perform the intended action via interaction with another 3P agent. Moreover, it may often be the case that the user is unaware of the availability of various 3P agents, and for various automated assistant interfaces it may be impractical and/or undesirable to explicitly provide a list of available 3P agents and associated functionality to a user in the often constrained automated assistant interface. For example, some automated assistants are “voice only” and it may be impractical and/or undesirable to “read a list” of 3P agents and associated functionality to the user.

Various implementations enable a user to engage in dialog with an automated assistant, and through that dialog the automated assistant may: determine an intended action of the user; determine value(s) for parameter(s) stored in association with the intended action; select a particular 3P agent that is able to achieve the intended action based on the determined values; and invoke the 3P agent with the determined values. Accordingly, instead of requiring the user start with an explicit invocation of a 3P agent, implementations disclosed herein instead engage in a dialog with a user in advance of invoking any 3P agent, and only invoke a particular 3P agent after determining the particular 3P agent can achieve the intended action with the determined values. These and other implementations may mitigate the waste of human and computational resources that may be encountered when a user explicitly invokes a 3P agent that turns out to be the wrong 3P agent. Moreover, in invoking the 3P agent, the automated assistant may transmit the determined values to the 3P agent and obviate the need for the 3P agent to solicit those values. This may enable the intended action